

Refine Search

Search Results -

Terms	Documents
L7 and (707/201).ccls.	2

Database:

US Pre-Grant Publication Full-Text Database
 US Patents Full-Text Database
 US OCR Full-Text Database
 EPO Abstracts Database
 JPO Abstracts Database
 Derwent World Patents Index
 IBM Technical Disclosure Bulletins

Search:

Refine Search

Recall Text
Clear
Interrupt

Search History

DATE: Wednesday, October 12, 2005 [Printable Copy](#) [Create Case](#)

<u>Set</u>	<u>Name</u>	<u>Query</u>	<u>Hit</u>	<u>Set</u>
			<u>Count</u>	<u>Name</u>
	<u>side by side</u>			<u>result set</u>
DB=USPT;		THES=ASSIGNEE; PLUR=YES; OP=OR		
<u>L11</u>		L7 and 707/201.ccls.	2	<u>L11</u>
<u>L10</u>		L7 and 707/200.ccls.	1	<u>L10</u>
<u>L9</u>		L7 and 707/10.ccls.	3	<u>L9</u>
<u>L8</u>		L7 and 707/1.ccls.	1	<u>L8</u>
<u>L7</u>		((logical near1 volume) same mirror) and master and slave	12	<u>L7</u>
<u>L6</u>		((master same slave) with (logical near1 volume))	0	<u>L6</u>
<u>L5</u>		L3 and ((logical same volume) and (master and slave))	0	<u>L5</u>
<u>L4</u>		L3 and (logical same volume)	2	<u>L4</u>
<u>L3</u>		L1 and (configura\$6 same client)	2	<u>L3</u>

<u>L2</u>	L1 and (configura\$6 saqme client)	8	<u>L2</u>
<u>L1</u>	(mirror\$4 with (file near1 system)) same virtual	9	<u>L1</u>

END OF SEARCH HISTORY


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used [mirror volume master slave logical configuration](#)

 Found **30,096** of **164,603**

Sort results by

[Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

[Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

[1 Fast detection of communication patterns in distributed executions](#)

Thomas Kunz, Michiel F. H. Seuren

 November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

 Full text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

[2 Highly available systems for database applications](#)

Won Kim

 March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1

 Full text available: [pdf\(2.43 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

[3 Hot mirroring: a method of hiding parity update penalty and degradation during rebuilds for RAID5](#)

Kazuhiko Mogi, Masaru Kitsuregawa

 June 1996 **ACM SIGMOD Record, Proceedings of the 1996 ACM SIGMOD international conference on Management of data**, Volume 25 Issue 2

 Full text available: [pdf\(1.37 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper proposes a storage management scheme for disk arrays, named hot mirroring. In this scheme, storage space is partitioned into two regions. One is the mirrored region, which is characterized by high performance and low storage efficiency. The other is the RAID5 region, which is characterized by low performance and high storage efficiency. Hot data blocks are stored in the former area, while cold blocks are stored in the latter. In addition,

mirrored pairs and RAID5 stripes are orthogona ...

4 Special issue on knowledge representation

Ronald J. Brachman, Brian C. Smith
February 1980 **ACM SIGART Bulletin**, Issue 70

Full text available:  pdf(13.13 MB) Additional Information: [full citation](#), [abstract](#)

In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt that there were two useful functions such an issue could serve. First, we hoped to elicit a clear picture of how people working in this subdiscipline understand knowledge representation research, to illuminate the issues on which current research is focused, and to catalogue what approaches and techniques are currently being developed. Second ...

5 An updated cross-indexed guide to the ray-tracing literature

L. Richard Speer
January 1992 **ACM SIGGRAPH Computer Graphics**, Volume 26 Issue 1

Full text available:  pdf(2.94 MB) Additional Information: [full citation](#), [index terms](#)

6 Doubly distorted mirrors

Cyril U. Orji, Jon A. Solworth
June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data**, Volume 22 Issue 2

Full text available:  pdf(1.05 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional mirrored disk systems provide high reliability by multiplexing disks. Performance is improved with parallel reads and shorter read seeks. However, writes must be performed by both disks, limiting performance. Doubly distorted mirrors increase the number of physical writes per logical write from 2 to 3, but performs logical writes more efficiently. This reduces the cost of a random logical write to 1/3 of the cost of a read. Moreover, much of the write ...

7 Illustrative risks to the public in the use of computer systems and related technology

Peter G. Neumann
January 1996 **ACM SIGSOFT Software Engineering Notes**, Volume 21 Issue 1

Full text available:  pdf(2.54 MB) Additional Information: [full citation](#)

8 AdaPT and Ada 9X

S. J. Goldsack, A. A. Holzbacher-Valero, R. Volz, R. Waldrop
March 1994 **ACM SIGAda Ada Letters**, Volume XIV Issue 2

Full text available:  pdf(1.18 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Following an initial proposal developed at the 4th Ada Real Time Programming Workshop, Nemacolin Woodlands, 1989, and further meetings to pursue the ideas, a series of papers were published [1][2][3] describing and discussing possible Ada language extensions to support the development of distributable programs. The language consisting of Ada with these extensions has been called AdaPT, which may be thought of as Ada with partitions. After introducing AdaPT, this paper describes how the AdaPT con ...

9

Evaluation of remote backup algorithms for transaction-processing systems

Christos A. Polyzois, Héctor García-Molina

September 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 3

Full text available:  [pdf\(1.87 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A remote backup is a copy of a primary database maintained at a geographically separate location and is used to increase data availability. Remote backup systems are typically log-based and can be classified into 2-safe and 1-safe, depending on whether transactions commit at both sites simultaneously or first commit at the primary and are later propagated to the backup. We have built an experimental database system on which we evaluated the performance of the epoch and the dependency recons ...

Keywords: disaster recovery, hot spare, hot standby, remote backup

10 The state of the art in locally distributed Web-server systems

Valeria Cardellini, Emiliano Casalicchio, Michele Colajanni, Philip S. Yu

June 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 2

Full text available:  [pdf\(1.41 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The overall increase in traffic on the World Wide Web is augmenting user-perceived response times from popular Web sites, especially in conjunction with special events. System platforms that do not replicate information content cannot provide the needed scalability to handle large traffic volumes and to match rapid and dramatic changes in the number of clients. The need to improve the performance of Web-based services has produced a variety of novel content delivery architectures. This article w ...

Keywords: Client/server, World Wide Web, cluster-based architectures, dispatching algorithms, distributed systems, load balancing, routing mechanisms

11 Evaluation of remote backup algorithms for transaction processing systems

Christos A. Polyzois, Hector Garcia-Molina

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data**, Volume 21 Issue 2

Full text available:  [pdf\(1.21 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A remote backup is a copy of a primary database maintained at a geographically separate location and is used to increase data availability. Remote backup systems are typically log-based and can be classified into 2-safe and 1-safe, depending on whether transactions commit at both sites simultaneously or they first commit at the primary and are later propagated to the backup. We have built an experimental database system on which we evaluated the performance of the epoch algorithm, a 1-safe ...

12 Novanet communications network for a control system

J. R. Hill, J. R. Severyn, P. J. VanArsdall

October 1983 **ACM SIGCOMM Computer Communication Review , Proceedings of the eighth symposium on Data communications**, Volume 13 Issue 4

Full text available:  [pdf\(1.07 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Novanet is a control system oriented fiber optic local area network that was designed to meet the unique and often conflicting requirements of the Nova laser control system which will begin operation in 1984. The computers and data acquisition devices that form the distributed control system for a large laser fusion research facility need reliable, high speed communications. Both control/status messages and experimental data must be handled. A subset of NOVANET is currently operating on the ...

13 A distributed UNIX system based on a virtual circuit switch

G. W.R. Luderer, H. Che, J. P. Haggerty, P. A. Kirslis, W. T. Marshall

December 1981 **Proceedings of the eighth ACM symposium on Operating systems principles**Full text available:  [pdf\(801.12 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The popular UNIXTM operating system provides time-sharing service on a single computer. This paper reports on the design and implementation of a distributed UNIX system. The new operating system consists of two components: the S-UNIX subsystem provides a complete UNIX process environment enhanced by access to remote files; the F-UNIX subsystem is specialized to offer remote file service. A system can be configured out of many computers which operate either under the S-U ...

14 FAB: building distributed enterprise disk arrays from commodity components

Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, Susan Spence

October 2004 **Proceedings of the 11th international conference on Architectural support for programming languages and operating systems**, Volume 32 , 38 , 39 Issue 5 , 5 , 11Full text available:  [pdf\(671.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the design, implementation, and evaluation of a Federated Array of Bricks (FAB), a distributed disk array that provides the reliability of traditional enterprise arrays with lower cost and better scalability. FAB is built from a collection of *bricks*, small storage appliances containing commodity disks, CPU, NVRAM, and network interface cards. FAB deploys a new majority-voting-based algorithm to replicate or erasure-code logical blocks across bricks and a reconfigurati ...

Keywords: consensus, disk array, erasure coding, replication, storage, voting

15 A distributed hypercube file system

R. J. Flynn, H. Hadimioglu

January 1989 **Proceedings of the third conference on Hypercube concurrent computers and applications - Volume 2**Full text available:  [pdf\(651.48 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

For the hypercube, an autonomous physically interconnected file system is proposed. The resulting distributed file system consists of an I/O organization and a software interface. The system is loosely-coupled architecturally but from operating systems point of view a tightly-coupled system is formed in which interprocessor messages are handled differently from file accesses. A matrix multiplication algorithm is given to show how the distributed file system is utilized.

16 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4Full text available:  [pdf\(5.33 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

17 Virtual reality learning environments: potentials and challenges

Meredith Bricken

July 1991 **ACM SIGGRAPH Computer Graphics**, Volume 25 Issue 3Full text available:  [pdf\(767.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This paper addresses the unique characteristics of emerging Virtual Reality (VR) technology and the potential of virtual worlds as learning environments. I describe several key attributes of VR environments and discuss them in relationship to educational theory and pedagogical practice. I then identify three challenges that must be met before VR can be integrated into educational settings: cost, usability, and fear of the technology.

18 The state of the art in distributed query processing

Donald Kossmann

December 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 4Full text available:  [pdf\(455.39 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the stat ...

Keywords: caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

19 Pipelined OR-parallelism architecture for parallel execution of Prolog

D. Sarma, C. P. Wu

June 1990 **Proceedings of the third international conference on Industrial and engineering applications of artificial intelligence and expert systems - Volume 2**Full text available:  [pdf\(1.14 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes a parallel architecture to implement PROLOG. The architecture considered here takes advantage of the OR-parallelism inherent in the language. In usual OR-parallelism architecture, multiprocessors are used to work on a relation of the database at the same time. In the pipeline OR-parallelism, a relation of the database is processed by only one of the multiprocessors at one time. As soon as a solution is found, the resolution goes forward and works on the goal down the go ...

20 Data communications control procedures for the USA standard code for information interchange

E. L. Lohse

March 1969 **Communications of the ACM**, Volume 12 Issue 3Full text available:  [pdf\(1.87 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This proposed American Standard has been accepted for publication by USA Standards Committee X3, Computers and Information Processing. In order that the final version of the proposed standard reflect the largest public consensus, X3 authorized publication of this document to elicit comment and general public reaction, with the understanding that such a working document is an intermediate result in the standardization process and is subject to change, modification, or withdrawal in p ...

Keywords: communication, communication control procedures, communication error control procedures, communication establishment/termination procedures, communication message transfer procedures, communication polling/selection procedures, data communication, data communication control procedures, data communication error control procedures, data communication establishment/termination procedures, data communication message transfer procedures, data communication polling/selection procedures, data link, data link control procedures, data link error control procedures, data link establishment, data link establishment/termination procedures, data link message transfer procedures, data link polling, data link polling/selection procedures, link, link control procedures, link error control procedures, link establishment/termination procedures, link message transfer procedures, link polling/selection procedures, selection procedures, termination procedures

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)